



HDD instability and the Rate of Failure

Hard Drive Failure

Google analysts have released a paper called "Failure Trends in a Large Disk Drive Population." This paper examines hard drive failure rates in Google's infrastructure. Two conclusions stood out:


1. Self-monitoring data isn't useful for predicting individual drive failures,
2. Temperature and activity levels don't correlate well with drive failure.

These comments have caused concern in the data recovery industry, so let us consider the findings in a little more detail.

The Google Study

The analysts examined the data from more than 100,000 drives deployed in Google's servers, all of which were consumer-grade serial and parallel ATA units with spindle speeds of 5400rpm and 7200rpm. Drives were considered "failed" if they were replaced as part of a repair procedure. Self-Monitoring, Analysis and Reporting Technology (SMART) information was recorded from all drives and spurious readings were filtered out of the resulting database.

When they looked at annualized failure rates, they saw the expected "infant mortality" effect, where drives die more often very early in their life cycle. The hypothesis behind this is that poorly-made drives fail quickly, while well-made ones have a few trouble-free years before they begin to reach their end-of-life stage which has been said to be around five years. This is sometimes referred to as the "bathtub curve" for its shape, but Google's researchers found that the failure rate ticked up much sooner—starting at two years—and remained steady for the next several years.



Data source:
Google



HDD instability and the Rate of Failure

Hard Drive Failure

When failures were correlated with SMART variables, the researchers found that only four SMART attributes had any relevance to the drive failure issue:

- Scan errors,
- Reallocation counts (when a drive remaps a bad sector to a "good" spare sector),
- Offline reallocations (a subset of the previous variable),
- Probational counts (the number of sectors "on probation" and suspected of being bad).

They also found that disk motors and spindles "don't fail". Google technicians examined the "spin retry" attribute, when a drive fails to spin-up the first time, and "did not register a single count within our entire population."

The four crucial SMART parameters are strongly correlated with higher drive failure levels, however the analysts were unable to use SMART information to create a meaningful predictive failure model that could warn engineers of impending drive failure. This was because 56 percent of all the failed drives in the study showed no counts in any of the four most important variables.

The Conclusion?

"SMART parameters still appear to be useful in analysing the aggregate reliability of large disk populations, which is still very important for logistics and supply-chain planning but not effective at predicting failure of an individual drive."

The researchers also found that drive failures did not increase with high temperatures or CPU utilisation. In fact, they say, lower average temperatures actually correlate more strongly with failure. Only at "very high temperatures" does this change.

Computer Science Labs Comment

Computer Science Labs Scientist Zijian Xie firstly considered the findings more relevant to hard disk drives that are deployed in specialist hosting facilities and server farms where drives are constantly active, and are in temperature-controlled environments. It is also very difficult to analyse the findings as there is no accurate reference to the temperature these drives were operating in.

Sponsored by the UK Government and the University of Manchester (a leading authority in the development of magnetic storage media), Zijian Xie continues to conduct his own intensive research into the use of non volatile magnetic media. Zijian Xie says:

"There is no doubt that an ambient operational temperature above that specified by the hard disk drive's manufacturer and temperature transitions affects the performance and longevity of magnetic storage media."

"Modern recording media is designed to be stable for at least 10 years at a disk operation temperature which is around 70 C". Adverse temperature rises causes a "thermal runaway effect across magnetic domains and also greater magnetic remanence of the media". As the drive is in use this problem manifests itself as Inter Symbol Interference (ISI) in the Digital Signal Processing System of a hard disk drive.



HDD instability and the Rate of Failure

Hard Drive Failure

Bit errors are consequently indicated in SMART and in sufficient magnitude will cause the hard disk drive to fail. This occurrence is common in many failed computer hard disk drive configurations and environments. Zijian Xie also considers the growing use of external hard disk enclosures a major contributing factor in hard disk storage failures.

The bigger issue, according to Zijian Xie, is the fact that drives fail too often. "Failure rates are always much higher than the manufacturers claim," he says, citing Google's data with an annual failure rate of eight percent for drives in service for two years. This equates to one out of every twelve drives failing! Computer Science Labs alongside University of Manchester is working to find sustainable magnetic storage methods and techniques.

Zijian Xie's observation is substantiated according to two Carnegie Mellon computer science professors. It has been found that drive failures are much higher in real-world situations than the manufacturer's mean time to failure (MTTF) number indicates.

"For drives less than five years old, field replacement rates were larger than what the datasheet MTTF suggested by a factor of 2-10," says the report. "For five to eight-year-old drives, field replacement rates were a factor of 30 higher than what the datasheet MTTF suggested." That's quite a difference between theory and practice, but it's not the only interesting information that was revealed by the study. The Carnegie Mellon group also found no "infant mortality" effect at all; instead, they saw a steadily-increasing rate of failures that began quite early in a drive's lifecycle. "This is an interesting observation," the researchers write, "because it does not agree with the common assumption that after the first year of operation, failure rates reach a steady state for a few years, forming the 'bottom of the bathtub'." The data was also gleaned from production enterprise systems and looked at a similar sample size to the Google study. It was not immediately clear why the two studies would produce such different results.